arXiv:1405.0524v3 [cs.GT] 18 May 2014

Computational Complexity of Approximate Nash Equilibrium in Large Games

Aviad Rubinstein*

May 20, 2014

Abstract

We prove that finding an ϵ-Nash equilibrium in a succinctly representable game with many players is PPAD-hard for constant ϵ. Our proof uses *succinct games*, i.e. games whose (exponentially-large) payoff function is represented by a smaller circuit. Our techniques build on a recent query complexity lower bound by Babichenko [1].

1 Introduction

Nash equilibrium is the central concept in Game Theory. Much of its importance and attractiveness comes from its *universality*: by Nash's Theorem [14], every finite game has at least one. The result that finding a Nash equilibrium is PPAD-complete, and therefore intractable [4, 2] casts this universality in doubt, since it suggests that there are games whose Nash equilibria, though existent, are in any practical sense inaccessible.

Can approximation repair this problem? Chen et al. [2] proved that it is also hard to find an ϵ-Nash equilibrium for any ϵ that is polynomially small - even for two-player games. The only remaining hope is a PTAS, i.e. an approximation scheme constant $\epsilon > 0$. *Whether there is a PTAS for the Nash equilibrium problem is the most important remaining open question in equilibrium computation.* When we say "ϵ-Nash equilibrium," we mean the additive sort; for multiplicative ϵ-Nash equilibria, Daskalakis [3] shows that the problem is PPAD-hard even for two-player games; notice that such a result is unlikely for additive approximation and a constant number of players, since a quasi-polynomial time approximation algorithm exists [13].

Our results

In this paper we make a modest step towards the inapproximability, in the standard additive sense, of Nash equilibrium:

*UC Berkeley. I am grateful to Christos Papadimitriou for inspiring discussions, comments, and advice. I would also like to thank Constantinos Daskalakis and Paul Goldberg for pointing out important missing references in an earlier version.

Theorem 1. *There exist constants $\epsilon, k > 0$, such that given a game with m players and k actions, and a poly-size circuit that computes the vector of payoffs for each vector of pure strategies, it is* PPAD*-hard to compute an ϵ-Nash equilibrium.*

Even though it is the first result establishing inapproximability of an additive notion of Nash equilibrium, it has two deficiencies: (a) it is about *ϵ-Nash-equilibrium* (sometimes[1] also called *ϵ-well-supported Nash equilibrium*) [4], requiring that all actions in the support be approximately best responses (instead of the mixture being best response in ϵ-approximate Nash equilibrium); and (b) it holds for a somewhat awkward class of multiplayer games we call *succinct games*[2], in which the utility of each player for an action profile is calculated by a circuit. For example, on such games computing the exact best response is PP-hard.

En route to proving Theorem 1, we also prove a similar statement about the computational complexity of finding an approximate fixed-point for a continuous (in fact, Lipschitz) function, which may be of separate interest.

Theorem 2. *There exist constants $\epsilon, M > 0$, such that given a M-Lipschitz, poly-time computable function $f\colon [0,1]^n \to [0,1]^n$, it is* PPAD*-hard to find an ϵ-approximate fixed point of f.*

Pure equilibria and PPAD-hardness It is interesting to note that all our hard game instances *have pure equilibria*. To the best of our knowledge, this is the first setting where PPAD-hardness is proved for games which have pure equilibria. Naturally, in any game with a poly-size payoff matrix, it is easy to find any pure equilibrium.

1.1 Related works

Schoenebeck and Vadhan [20] studied comprehensively the computational complexity of Nash Equilibria in succinct games and other concise representations. They characterized the computational complexity, for different levels of approximation, of the following questions: is a given strategy profile a Nash equilibium? does there exist a pure equilibrium? does there exist a Nash equilibrium in which the payoff to a given player meets a given guarantee? They also prove that the problem of finding an ϵ-Nash equilibrium (for constant ϵ) in circuit games where each player has two strategies is BPP-hard (for a promise-problem version of BPP), and belongs to P^{MA}. Finding a tighter complexity classification is stated as an open question.

We note that different variants of succinct games were also studied by [18, 7, 8].

[1] See [3] for a short discussion on terminology.

[2] There is some disagreement in the literature about the terminology: a very similar definition appeared in [20] as *circuit games*; [7] discuss *polynomially definable games*; while [8] did use the term succinct games.

Related works on query complexity

There are several interesting results on the query complexity of approximate Nash equilibria, where the algorithm is assumed to have black-box access to the exponential-size payoff function. In other words, in this setting the payoff function is allowed to be arbitrarily complex.

A recent paper by Hart and Nisan [11] proves that any deterministic algorithm needs to query at least an exponential number of queries to compute any ϵ-Nash equilibrium - and even for any ϵ-correlated equilibrium. For ϵ-correlated equilibrium, on the other hand, Hart and Nisan show a randomized algorithm that uses a number of queries polynomial in n and ϵ^{-1}.

Of particular interest to us, is a very recent paper by Babichenko [1], that extends the hardness result of Hart and Nisan to show that any randomized algorithm still requires an exponential number of queries to find an ϵ-Nash equilibrium. Our proof is inspired by Babichenko's work and builds on his techniques.

Finally, yet a newer paper by Goldberg and Roth [10] characterizes the query complexity of approximate coarse correlated equilibrium in games with many players. More important for our purpose is their polynomial upper bound on the query complexity of ϵ-Nash equilibria for any family of games that have *any* concise representation. This result is to be contrasted with (a) Babichenko's query complexity lower bound, which uses a larger family of games, and (b) our result which applies exactly to this setting and gives a lower bound on the *computational complexity*.

A much older yet very interesting and closely related result is that of Hirsch, Papadimitriou, and Vavasis [12]. Hirsch et al. show that any deterministic algorithm for computing a Brouwer fixed point in the oracle model must make an exponential -in the dimension n and the approximation ϵ- number of queries for values of the function. The techniques in [12] have proven particularly useful both in Babichenko's work, and in ours.

Related works on succinct representations of other objects

Although our notion of succinct representation is somewhat non-standard in game theory, similar succinct problems have been considered before. Galperin and Wigderson [9] and Papadimitriou and Yannakakis [17] studied the computation of graph properties for exponential-size graphs given access to a circuit that locally computes the adjacency matrix.

Much more recently, Dobzinski and Vondrak studied optimization of succinctly-represented objects for submodular optimization [6] and for combinatorial auctions [5]. As in our case, in both of those settings similar hardness results were previously known in the value oracle model.

1.2 Techniques

Our techniques in this paper are significantly different from previous works on PPAD-hardness in games such as Daskalakis et al. [4] and Chen et al. [2]. In

particular two ingredients of our reduction are fundamentally different from the so-called "DGP framework": the construction of the Brouwer function, and the reduction from Brouwer to Nash. We note that in both cases our techniques are much closer to the recent work of Babichenko [1].

From Brouwer to Nash Our reduction from Brouwer to Nash follows an argument presented in Eran Shmaya's blog [21]. Each player has actions in $[0, 1]$, and the players are divided into two groups: the first group tries to imitate the second group, while the second group tries to imitate the Brouwer function applied to the actions chosen by the first group.

In [4, 2], per contra, the probability assigned to each action corresponds, roughly, to a variable in $[0, 1]$. While this construction is more stringent on the number of actions, it requires a relatively complex averaging gadget. More importantly, this averaging gadget as implemented by Chen et al. seems to require a polynomial blow-up of the error. Thus, it is not clear how to achieve a constant hardness of approximation in this way.

The Brouwer function In order to construct the hard Brouwer function we use a construction due to Hirsch, Papadimitriou, and Vavasis [12] that embeds paths as mappings over $[0, 1]^n$, in a delicate way which we describe later.

Chen et al. [2], on the other hand, divide the unit hypercube into subcubes of constant edge-length, and specify a color for each subcube. The i-th color corresponds to ξ_i, i.e. the i-th unit vector, whereas a special red color corresponds to $-\sum \xi_i$. Any vector of players' mixed strategies corresponds to a distribution over neighboring subcubes, and we get a fixed point whenever the expectation of the corresponding vectors is 0^n.

Clearly, any distribution corresponding to an exact fixed point must have support over a panchromatic neighborhood, i.e. a neighborhood with subcubes corresponding to all $n + 1$ colors. Finally, it is shown that it is PPAD-hard to find such panchromatic neighborhood. In fact, a panchromatic neighborhood is still necessary for any $\Theta(1/n)$-approximate[3] fixed point. However, a $(1/n)$-approximate fixed point can be achieved from a distribution over only n colors: for example, taking each of the colors $i \in [n]$ with probability $1/n$ results in the expected vector $(1/n, \ldots, 1/n)$.

2 Preliminaries

Throughout this paper we use the max-norm as the default measure of distance. In particular, when we say that f is M-Lipschitz we mean that for every $\mathbf{x}$ and $\mathbf{y}$ in the domain of f, $\|f(\mathbf{x}) - f(\mathbf{y})\|_\infty \leq M \|\mathbf{x} - \mathbf{y}\|_\infty$.

[3] Approximation here is in the sense of L^∞. In fact, it seems to give a constant hardness of approximation in L^1, but unfortunately it is not clear how to use that for a reduction to Nash equilibria.

The EndOfTheLine problem

Our reduction starts form the EndOfTheLine problem. This problem was implicit in [16], and explicitly by defined Daskalakis et al. [4].

Definition 1. EndOfTheLine: ([4]) Given two circuits S and P, with n input bits each, such that $P(0^n) = 0^n \neq S(0^n)$, find an input $x \in \{0,1\}^n$ such that $P(S(x)) \neq x$ or $S(P(x)) \neq x \neq 0^n$.

Theorem. *(Essentially [16])* EndOfTheLine *is* PPAD*-complete (for poly-size S and P).*

Succinct normal form games

We consider games with n players. Player i chooses one of k actions $a_i \in A_i$. The utility of player i for each vector of actions is given by $u_i \colon \times_j A_j \to [0,1]$.

Explicitly describing the u_i's requires space exponential in n. In this paper we restrict our attention to games that can be described *succinctly*; namely, there is a poly-size circuit that computes each of the u_i's given a vector of actions $\mathbf{a} \in \times_j A_j$.

ϵ-Nash equilibrium vs ϵ-approximate Nash equilibrium

A mixed strategy of player i is a distribution $x_i \in \Delta A_i$. We say that a vector of mixed strategies $\mathbf{x} \in \times_j \Delta A_j$ is a *Nash equilibrium* if every strategy a_i in the support of x_i is a best response to the actions of the mixed strategies of the rest of the players, x_{-i}. Formally,

$$\mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a_i, a_{-i})] = \max_{a' \in A_i} \mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a', a_{-i})] \,.$$

Equivalently, $\mathbf{x}$ is a Nash equilibrium if each mixed strategy x_i is a best mixed response to x_{-i}:

$$\mathbb{E}_{\mathbf{a} \sim \mathbf{x}} [u_i(\mathbf{a})] = \max_{x_i' \in \Delta A_i} \mathbb{E}_{\mathbf{a} \sim \mathbf{x}'} [u_i(\mathbf{a})] \,.$$

Each of those equivalent definitions can be generalized to include approximation in a different way. (Of course, there are also other interesting generalizations of Nash equilibria to approximate settings.) We say that $\mathbf{x}$ is an ϵ*-approximate Nash equilibrium* (ϵ*-ANE*) if each x_i is an ϵ-best mixed response to x_{-i}:

$$\mathbb{E}_{\mathbf{a} \sim \mathbf{x}} [u_i(\mathbf{a})] \geq \max_{x_i' \in \Delta A_i} \mathbb{E}_{\mathbf{a} \sim \mathbf{x}'} [u_i(\mathbf{a})] - \epsilon \,.$$

On the other hand, we generalize the first definition of Nash equilibrium by saying that $\mathbf{x}$ is a ϵ*-Nash equilibrium* (ϵ*-NE*; sometimes also ϵ-well-supported Nash equilibrium) if each a_i in the support of x_i is an ϵ-best response to x_{-i}:

$$\mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a_i, a_{-i})] \geq \max_{a' \in A_i} \mathbb{E}_{a_{-i} \sim x_{-i}} [u_i(a', a_{-i})] - \epsilon \,.$$

It is easy to see every ϵ-NE is also an ϵ-ANE, but the converse is false. Given an ϵ-ANE it is possible to find a $\Theta(\sqrt{\epsilon} n)$-NE (see e.g. [4]); however computational hardness of n^{-c}-approximate Nash equilibrium is a corollary of [2].

3 Main result

Definition 2. Given a game over m players with k actions, and a poly-size circuit that computes the vector of payoffs for each vector of pure strategies, SuccinctNash(n, k, ϵ) is the problem of computing an ϵ-NE for this game.

Theorem 1. SuccinctNash$\left(n, 10^4, 10^{-8}\right)$ is PPAD-hard.

Proof overview

We begin our proof with the EndOfTheLine problem on $\{0,1\}^n$ [4]. In the first step, we embed the EndOfTheLine as a collection H of vertex-disjoint paths on the $(2n+1)$-dimensional hypercube graph. Given H, our second step is to construct a continuous mapping $f\colon [0,1]^{2n+2} \to [0,1]^{2n+2}$ whose fixed points correspond to ends of paths in H. This step is done using a technique introduced by Hirsch et al. [12]. Our third and final step is to reduce the problem of finding approximate fixed points of f to the problem of finding approximate NE in a $4n+4$-players game via a reduction which appeared in Shmaya's blog [21].

3.1 Embedding the EndOfTheLine graph as paths in $\{0,1\}^{2n+1}$

Our first step in the reduction is to embed an EndOfTheLine graph as vertex-disjoint paths on the $(2n+1)$-dimensional hypercube graph. We first recall that the input to the EndOfTheLine problem is given as two circuits S and P, which define a directed graph over G over $\{0,1\}^n$. Given S and P, we construct a collection H of vertex-disjoint paths over the $(2n+1)$-dimensional hypercube graph, such that each starting or end point of a path in H corresponds to a unique starting or end point of a line in G.

In order to construct our embedding we divide the $2n+1$ coordinates as follows: the first n coordinates store current vertex $\mathbf{u}$, the next n coordinates for next vertex in the line, $\mathbf{v}$, and finally, the last coordinate b stores a compute-next vs copy bit. When $b=0$, the path proceeds to update $\mathbf{v} \leftarrow S(\mathbf{u})$, bit-by-bit. When this update is complete, the value of b is changed to 1. Whenever $b=1$, the path proceeds by copying $\mathbf{u} \leftarrow \mathbf{v}$ bit-by-bit, and then changes that value of b again. Finally, when $\mathbf{u} = \mathbf{v} = S(\mathbf{u})$ and $b=0$, the path reaches an end point.

Notice that the paths in H do not intersect. Furthermore, given a vector in $\mathbf{p} \in \{0,1\}^{2n+1}$, we can output in polynomial time whether $\mathbf{p}$ belongs to a path in H, and if so which are the previous and consecutive vectors in the path. It is therefore PPAD-hard to find a starting or end point of any path in H other than 0^n.

3.2 Continuous mapping on $[0,1]^{2n+2}$

Our second step, which constructs a continuous mapping given H, is probably the most technically involved. Fortunately, almost all of the technical work we need was already done by Hirsch et al. [12].

Definition 3. Given an M-Lipschitz, poly-time computable function $f\colon [0,1]^n \to [0,1]^n$, SuccinctBrouwer$(n, M, \epsilon)$ is the problem of computing an ϵ-approximate fixed point of f.

Theorem 2. SuccinctBrouwer$(n, 80, 1/88)$ is PPAD-hard.

Proof. We begin with a quick overview of the mapping constructed by Hirsch et al. [12]. We will then show how to adapt their construction to fit our reduction. In the following, we denote $g(\mathbf{x}) = f(\mathbf{x}) - \mathbf{x}$.

The HPV mapping

Given a path in $\{0,1\}^{2n+1}$, Hirsch et al. [12] construct a mapping $f\colon [0,1]^{2n+2} \to [0,1]^{2n+2}$ that satisfies[4]:

1. g is 79-Lipschitz (thus, f is 80-Lipschitz)
2. $\|g(\mathbf{x})\|_\infty \geq 1/88$ for every $\mathbf{x}$ that does not correspond to the endpoint of the path
3. The value of g at each point $\mathbf{x}$ depends only on whether the path passes through the subcube corresponding to $\mathbf{x}$, and in which direction.

However, for our purposes, it does not suffice to embed just a single path. In particular, we have one path whose starting point we know (and must not correspond to a fixed point), and many other paths (and cycles), whose end points -and starting points- correspond to additional fixed points.

Luckily, the same construction of [12] continues to work in our case, with only minor changes. In order to explain those modifications, we first briefly recall some of the details in the original construction.

We divide the $[0,1]^{2n+2}$ hypercube into smaller subcubes of edge-size h. In [12], a single path over $\{0,1\}^{2n+1}$ corresponds to a $(2n+1)$-dimensional sequence of subcubes, called the *tube*, that all lie in a special designated *slice* of $[0,1]^{2n+2}$.

The *home subcube*, the subcube that corresponds to the beginning of the path, is special: the flow from all subcubes that do not belong to the path leads to this subcube.

For the purpose of adapting this construction, the most important property is that on (and near) the outer facets of the tube, i.e. the facets that do not face other subcubes in the path, $g(\mathbf{x}) = \delta\xi_{2n+2}$, where ξ_{2n+2} is the $(2n+2)$-unit vector, and δ is some small parameter (constant in this paper). The same also holds for all points in the slice that do not belong to the tube. Intuitively, this means that *all subcubes -whether they belong to the tube or not- look the same from the outside* (except for the two facets that continue the path).

[4] See also [1] for the choice of constants. Please note the change in variable names: $h, M, 2^{-p}$ in [12] correspond to $\delta, \lambda, \epsilon$, respectively, in [1].

Embedding multiple paths

Given a collection H of non-intersecting paths in $\{0,1\}^{2n+1}$, we construct a mapping $f\colon [0,1]^{2n+2} \to [0,1]^{2n+2}$ in a similar fashion. Essentially, we construct many tubes (some of which may form close cycles) in the same slice.

If any path in H passes through some $\mathbf{p} \in \{0,1\}^{2n+1}$, then $\mathbf{p}$ corresponds to a subcube on which f is defined exactly the same as in [12]. Likewise, every end point of path in H corresponds to a subcube on which f is defined exactly the same as the unique end point in [12].

Our construction differs in the starting points of the paths in H. In [12], the starting point corresponds to the home subcube, which is *universally unique*: the flow from any subcube outside the path leads towards that point. Indeed we cannot imitate this behaviour for the starting point of every path in H. However, this does not complicate our reduction, because all other starting points are also (PPAD-) hard to find. In particular we can construct f in a similar fashion to the end points of the paths, thereby creating additional fixed points.

For each starting point, consider the corresponding subcube of $[0,1]^{2n+2}$. The values of f on the facet F_1 in the direction of the path, are already determined by the next subcube in the path. Now, we let $g(\mathbf{x}) = \delta\xi_{2n+2}$ uniformly for every point $\mathbf{x}$ on the opposite facet, F_0. For any $\mathbf{x}'$ between the facets we interpolate by taking the weighted average of the values $f\left(\mathbf{x}_0^{'}\right)$ and $f\left(\mathbf{x}_1^{'}\right)$ on the projections of $\mathbf{x}'$ on F_0 and F_1, respectively. (This corresponds to [12]'s "Cartesian interpolation"). Notice that this subcube also satisfies $g(\mathbf{x}) = \delta\xi_{2n+2}$ for all points on facets that do not face the rest of the path.

Outside the subcubes corresponding to the additional paths, all the properties of the mapping are trivially preserved. On the interface between any subcube in an additional path, and any other subcube which is not consecutive in the same path, all the properties are again preserved since adding the paths does not change the value of f near those facets. Finally within each path, f is constructed exactly the same way as the single path in [12], and therefore it is easy to see that all the properties continue to hold - except of course the new fixed points at the starting points of the additional paths.

Finally, we conclude that it is PPAD-hard to find an approximate fixed point of f. □

3.3 From Brouwer to WSNE

Our third and final step in the proof reduces the problem of finding approximate fixed points in f to that of finding approximate well-supported Nash equilibria. The reduction we use is based on [21] and appears almost exactly in this format in [1].

Proposition 1. *(Essentially [1])* $\textsc{SuccinctBrouwer}(n, M, \epsilon) \leq_P \textsc{SuccinctNash}\left(2n, k+1, \frac{3}{4k^2}\right)$, *where* $k = \lceil \frac{3+M}{\epsilon} \rceil$

Proof. We construct a game with two groups of $2n+2$ players each. The action set of each player corresponds to $\{0, 1/k, \ldots, 1\}$. We denote the choice of

strategies for the first group $\mathbf{a} = (a_1 \dots a_{2n+2})$, and $\mathbf{b} = (b_1, \dots b_{2n+2})$ for the second group.

Each player in the first group attempts to imitate the behaviour of the corresponding player in the second group. Her utility is given by

$$u_i(a_i, b_i) = -|a_i - b_i|^2$$

The second group players attempt to imitate the value of f, when applied to the vector of actions taken by all the players in the first group. The utility of the i-th player is

$$v_i(b_i, \mathbf{a}) = -|f_i(\mathbf{a}) - b_i|^2$$

Observe that when the i-th player on the second group (henceforth, player $(i, 2)$) applies a mixed strategy, the expected utility for player $(i, 1)$ (the i-th player in the first group) is given by:

$$\mathbb{E}[u_i(a_i, b_i)] = -|a_i - \mathbb{E}(b_i)|^2 - \mathbf{Var}(b_i)$$

Let $\alpha_i \in \{0, 1/k, \dots, 1\}$ be such that $\mathbb{E}(b_i) \in [\alpha_i, \alpha_i + 1/k]$, and assume wlog that $\mathbb{E}(b_i) \in [\alpha_i, \alpha_i + 1/2k]$. Then we can lower-bound the expected utility of player $(i, 1)$ when playing α_i:

$$\mathbb{E}[u_i(\alpha_i, b_i)] \geq -\frac{1}{4k^2} - \mathbf{Var}(b_i)$$

On the other hand, for any $\gamma \notin \{\alpha_i, \alpha_i + 1/k\}$,

$$\mathbb{E}[u_i(\gamma, b_i)] \leq -\frac{1}{k^2} - \mathbf{Var}(b_i)$$

Therefore in every $\frac{3}{4k^2}$-NE, the support of the mixed strategy of player $(i, 1)$ is restricted to $\{\alpha_i, \alpha_i + 1/k\}$.

For the second group of players, we have

$$\mathbb{E}[v_i(b_i, \mathbf{a})] = -|\mathbb{E}(f_i(\mathbf{a})) - b_i|^2 - \mathbf{Var}(f_i(\mathbf{a}))$$

Let $\beta_i \in \{0, 1/k, \dots, 1\}$ be such that $\mathbb{E}(f_i(\mathbf{a})) \in [\beta_i, \beta_i + 1/k]$, then in every $\frac{3}{4k^2}$-NE, the support player $(i, 2)$ is restricted to $\{\beta_i, \beta_i + 1/k\}$.

Finally, given a $\frac{3}{4k^2}$-NE, we use the Lipschitz property of f to derive an approximate fixed point. Notice that for α and β as defined above,

$$|\alpha_i - \beta_i| \leq |\alpha_i - \mathbb{E}(b_i)| + |\mathbb{E}(b_i) - \beta_i| \leq \frac{2}{k}$$

Likewise,

$$|\beta_i - f_i(\alpha)| \leq |\beta_i - \mathbb{E}(f_i(\mathbf{a}))| + |\mathbb{E}(f_i(\mathbf{a})) - f_i(\alpha)| \leq \frac{1}{k} + \frac{M}{k},$$

where M is the Lipschitz constant of f.

Therefore $|\alpha - f(\alpha)|_\infty \leq \frac{3+M}{k}$, so α is a $\frac{3+M}{k}$-approximate fixed point of f. □

4 Open problems

We prove that finding an ϵ-well-supported Nash equilibrium PPAD-hard even for constant ϵ. This is a modest step towards a better understanding of the computational complexity of approximating Nash equilibrium in games with a large number of players. Many important questions remain open:

NE vs WSNE As mentioned earlier, in the domain of constant approximations, Nash equilibria may be strictly easier than finding well-supported Nash equilibria.

What is the complexity of finding ϵ-ANE?

We note that a similar obstacle was encountered by Babichenko [1], in the case of query complexity, for essentially the same reasons.

Simpler games Other, more restrictive forms of succinct games have been studied quite extensively before (see e.g. [15]). It would be interesting to improve our understanding of the complexity of those games. In particular,

What is the complexity of finding an ϵ-NE in bounded-degree graphical games?

Upper bound on complexity Embarrassingly, we do not have a (non-trivial) upper bound on the complexity of finding an ϵ-NE. One particular obstacle to showing that this problem belongs to the class PPAD, is that it is not even clear how to find an approximate best response to mixed strategies without using randomness. (In particular, by [20], this problem is also BPP-hard; we do not know whether BPP is contained in PPAD.)

What is the right complexity classification of the problem of finding an ϵ-NE?

5 Addendum

Since the first version of this paper was posted, we found solutions to the first two open problems mentioned above [19].

References

[1] Yakov Babichenko. Query complexity of approximate nash equilibria. *CoRR*, abs/1306.6686, 2013. To appear in STOC 2014.

[2] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *J. ACM*, 56(3), 2009.

[3] Constantinos Daskalakis. On the complexity of approximating a nash equilibrium. *ACM Transactions on Algorithms*, 9(3):23, 2013.

[4] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. *Commun. ACM*, 52(2):89–97, 2009.

[5] Shahar Dobzinski and Jan Vondrák. The computational complexity of truthfulness in combinatorial auctions. In *EC*, pages 405–422, 2012.

[6] Shahar Dobzinski and Jan Vondrák. From query complexity to computational complexity. In *STOC*, pages 1107–1116, 2012.

[7] Joan Feigenbaum, Daphne Koller, and Peter W. Shor. A game-theoretic classification of interactive complexity classes. In *Structure in Complexity Theory Conference*, pages 227–237, 1995.

[8] Lance Fortnow, Russell Impagliazzo, Valentine Kabanets, and Christopher Umans. On the complexity of succinct zero-sum games. *Computational Complexity*, 17(3):353–376, 2008.

[9] Hana Galperin and Avi Wigderson. Succinct representations of graphs. *Inf. Control*, 56(3):183–198, April 1984.

[10] Paul W. Goldberg and Aaron Roth. Bounds for the query complexity of approximate equilibria. In *EC*, 2014. To appear.

[11] Sergiu Hart and Noam Nisan. The query complexity of correlated equilibria. *CoRR*, abs/1305.4874, 2013.

[12] Michael D. Hirsch, Christos H. Papadimitriou, and Stephen A. Vavasis. Exponential lower bounds for finding brouwer fix points. *J. Complexity*, 5(4):379–416, 1989.

[13] Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *EC*, pages 36–41, 2003.

[14] John Nash. Non-cooperative games. *The Annals of Mathematics*, 54:286–295, 1951.

[15] Chistos Papadimitriou. *Succint Representations of Games*, chapter 2.5, pages 39–41. Cambridge University Press, 2007.

[16] Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994.

[17] Christos H. Papadimitriou and Mihalis Yannakakis. A note on succinct representations of graphs. *Information and Control*, 71(3):181 – 185, 1986.

[18] Christos H. Papadimitriou and Mihalis Yannakakis. On complexity as bounded rationality (extended abstract). In *STOC*, pages 726–733, 1994.

[19] Aviad Rubinstein. Inapproximability of nash equilibrium. *CoRR*, abs/1405.3322, 2014.

[20] Grant Schoenebeck and Salil P. Vadhan. The computational complexity of nash equilibria in concisely represented games. *TOCT*, 4(2):4, 2012.

[21] Eran Shmaya. Brouwer implies nash implies brouwer. http://theoryclass.wordpress.com/2012/01/05/brouwer-implies-nash-implies-brouwer/, 2012.